

LONG-TERM CLOCK BEHAVIOR OF GPS IIR SATELLITES

Dr. Marvin Epstein, Todd Dass, Dr. John Rajan, and Paul Gilmour

E-mail: *marvin.epstein@itt.com, todd.dass@itt.com,*

john.rajan@itt.com, paul.gilmour@itt.com

Tel: 973.284.2614, 719.567.3928, 973.284.4835, 973.284.2523

ITT Corporation

Space Systems Division

77 River Road, Clifton, NJ 07014, USA

Abstract

The ITT-developed GPS IIR satellite payloads have been on orbit since 1997, and have proven to be the best family of clocks in the GPS constellation. At this time, there is a substantial recorded history of clock behavior, including over 60 clock years of space operation. The age of the oldest clock is over 9 years. A review of the record shows a number of significant characteristics that were not apparent in shorter clock tests. Rubidium clocks, as opposed to cesium clocks, have significant long-term drift. The current literature describes an initial model of drift aging for rubidium atomic clocks followed by a long-term characteristic. A review of the IIR clocks shows what appears to be another significant break point in the long-term drift characteristics. The usual assumption is that the drift tends toward zero drift in the long term. The data indicate that the long-term drift will always remain negative and stay substantially away from zero. It is commonly known that some rubidium clocks generate frequency steps and the frequency steps tend to decrease in size and rate of occurrence over time. We have seen a number of cases of this behavior. Most frequency steps tend to be frequency steps that tend to persist in time, but we also have seen triangular frequency patterns where a sudden jump in frequency is followed by a rapid decay to the long-term drift pattern. Also, we have cases in which a frequency step pattern that has persisted for a significant time suddenly disappears. Although there are cases where the frequency step intensity initially grows after turn-on, the overall pattern is for the frequency step intensity to decay over the long term. Future observations of these clocks will allow us to determine if this decay in frequency step intensity is permanent.

I. INTRODUCTION

The GPS Block IIR clocks are the best performing and largest family of clocks in the current GPS constellation. We have collected a large amount of data about these clocks, which offers a unique opportunity for an analysis of the behavior of rubidium clocks. As of 13 August 2006, we had 56 years of life in orbit for these clocks. There are 13 satellites in orbit with these clocks and the ages of the operational clocks range from 0.9 years to 9 years. Seven more satellites are on the ground and when they are launched, there will be a total of 20 such satellites.

Report Documentation Page			Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.				
1. REPORT DATE NOV 2007		2. REPORT TYPE		3. DATES COVERED 00-00-2007 to 00-00-2007
4. TITLE AND SUBTITLE Long-Term Clock Behavior of GPS IIR Satellites		5a. CONTRACT NUMBER		
		5b. GRANT NUMBER		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)		5d. PROJECT NUMBER		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) ITT Corporation,Space Systems Division,77 River Road,Clifton,NJ,07014		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)		
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited				
13. SUPPLEMENTARY NOTES 39th Annual Precise Time and Time Interval (PTTI) Meeting, 26-29 Nov 2007, Long Beach, CA				
14. ABSTRACT The ITT-developed GPS IIR satellite payloads have been on orbit since 1997, and have proven to be the best family of clocks in the GPS constellation. At this time, there is a substantial recorded history of clock behavior, including over 60 clock years of space operation. The age of the oldest clock is over 9 years. A review of the record shows a number of significant characteristics that were not apparent in shorter clock tests. Rubidium clocks, as opposed to cesium clocks, have significant long-term drift. The current literature describes an initial model of drift aging for rubidium atomic clocks followed by a long-term characteristic. A review of the IIR clocks shows what appears to be another significant break point in the long-term drift characteristics. The usual assumption is that the drift tends toward zero drift in the long term. The data indicate that the long-term drift will always remain negative and stay substantially away from zero. It is commonly known that some rubidium clocks generate frequency steps and the frequency steps tend to decrease in size and rate of occurrence over time. We have seen a number of cases of this behavior. Most frequency steps tend to be frequency steps that tend to persist in time, but we also have seen triangular frequency patterns where a sudden jump in frequency is followed by a rapid decay to the long-term drift pattern. Also, we have cases in which a frequency step pattern that has persisted for a significant time suddenly disappears. Although there are cases where the frequency step intensity initially grows after turn-on, the overall pattern is for the frequency step intensity to decay over the long term. Future observations of these clocks will allow us to determine if this decay in frequency step intensity is permanent.				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 20
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified		

Typical ground-based tests of atomic clocks are shorter and involve fewer clocks. The reason we have such a large body of data is that these clocks are part of the GPS constellation and, as such, the vacuum comes for free and we automatically have personnel to manage the clocks, collect the data, and document them. Also, these clocks use a single design and were produced in one run with deliveries over a 4.5-year interval (December 1994 to May 1999). This allows us to examine the variation in behavior of these clocks, without being involved with complications of different technologies, designs, and production runs. Each satellite has three clocks onboard so that there will be 60 of these clocks in the GPS constellation. Standard operation is to have one clock powered up and serving as the operational clock. The other two clocks are unpowered and serve as backup. Because of the unique nature of this body of data, one can't simply generalize the results of these data to other rubidium clock designs.

There are a few points that we want to make about our notation. We will refer to these clocks, in the following, by the acronym RAFS (Rubidium Atomic Frequency Standard), to simplify the discussion. Secondly, the Block IIR RAFS and the associated TKS (Time-Keeping System) [1] are used on all 20 satellites, but there are 12 satellites that are labeled as Block IIR satellites and 8 satellites as Block IIRM satellites. The Block IIR and IIRM satellites use the same Block IIR clock design, but the Block IIRM satellites have more transmitter power capability than do the Block IIR satellites. Also, the Block IIRM satellites support both the heritage GPS waveforms and some modernized waveforms, while the Block IIR satellites only support the heritage power level and waveforms. Satellites from the Block IIR/IIRM family before 1 January 2005 are Block IIR satellites. Satellites launched after this date are Block IIRM satellites. To emphasize the fact that the RAFS is the same design on both Block IIR and IIRM satellites, we will refer to these clocks throughout the paper as Block IIR clocks. We will refer to the satellites by their SVN (Satellite Vehicle Number), as this is the nomenclature used by the Air Force. Up to now, each of the IIR/IIRM satellites is on its original operational clock. Thus, in the following, we will generally refer to an operational clock by its associated SVN number. Technically, we should reference the clock by its SVN number and its clock number (e.g., clock 1, 2, or 3).

This paper will analyze the long-term drift and large frequency step characteristics observed in our RAFS data. The following sections describe the reasons for this selection, the processing of the received clock signal needed to obtain the drift and frequency step data, and our analysis of these effects.

II. SELECTION OF STUDY TOPICS

The basic GPS analysis starts with the observed pseudo-ranges at various monitor stations. The pseudo-ranges are processed by various organizations to produce an estimate of satellite position, velocity and time, using a Kalman filter or equivalent. Our analysis starts with the satellite clock phase and frequency as reported by NGA (the National Geospatial-Intelligence Agency) [2]. We have selected the NGA data for our analysis because it covers the full range of the GPS Block IIR life and it provides both phase and frequency estimates.

NGA clock data are the result of many inputs such as TKS effects and RAFS effects with both short-term and long-term components. Thus, these data are the basis for a variety of interesting potential study topics. We decided to concentrate on RAFS effects for this paper since the RAFS is the key source of local reference time on the satellite. The first topic we picked to study was the drift of the RAFS. Drift is a well known, significant characteristic of rubidium atomic clocks. This is in contrast to cesium clocks, whose drift is very close to zero. Figure 1 shows that the range of frequencies for the operational IIR clocks is about $1 \cdot 10^{-9}$. This figure also provides the age of the operational IIR/IIRM satellites and the frequency range that they have spanned. Figure 2 shows the frequency range and clock age for each satellite. It shows that, while in general the clocks that are older have drifted further, this is not always

the case (see SVN46). Figure 3 shows the drift patterns of all the operational IIR clocks using a 100-day frequency difference to estimate the drift. We see that IIR clocks have a similar long-term drift pattern. We hypothesized that we could make a good mathematical approximation to the drifts of the various IIR clocks. This paper describes the results of our analysis.

Note that drift is a derivative of the frequency and the noisiness of the drift depends on the length of time over which the drift is averaged. Short-term drift tends to be dominated by short-term noise. Therefore, we concentrated our analysis on the long-term drift (such as drift averaged over a span longer than 14 days). This longer averaging results in a higher signal-to-noise ratio on the drift estimate.

The second topic we picked was frequency steps. Frequency steps come in a variety of sizes ranging from very small to large. We decided to concentrate on the large frequency steps because: a) they are the most significant anomaly of this family of clocks and heavily affect their Hadamard deviation, b) they are easily distinguishable from the noise, and c) they are fairly small in number, so enumerating and processing these jumps is manageable. We have no magic threshold for differentiating between large and small frequency steps, as we assume that the probability distribution of frequency step size is continuous. For our analysis, we hand-selected the large frequency steps looking for steps that generally were larger than $1 \cdot 10^{-13}$ in size and visibly affected the frequency for over 5 days.

There is synergy in the study of long-term drift and large frequency steps: once most of the drift is removed, the large frequency steps become much more visible. On the other hand, once the large frequency steps are removed, one can make a better fit to the drift.

III. PROCESSING OF THE NGA FREQUENCY

The NGA data phase and frequency values are spaced 15 minutes apart, with the first data point in an hour at 00 minutes and 00.0 seconds. This means that there are roughly 315,600 points in 9 years, which is more data than we could conveniently handle. Thus, as shown in the top block of Figure 4, we thinned the data out by selecting one point out of every day. The NGA data consist of a sequence of 1-day fits, where the fit for a given target day is based on the monitor station data for the target day, the day before the target day, and the day following the target day. Each day, we selected the data point in the middle of the fit, namely the noon data point, as our sample value.

The NGA data represent the value of the broadcast clock, which shows the effects of both the RAFS and the TKS system. Figure 5 shows the broadcast clock frequency for SVN41 as a function of days after turn-on. SVN41 is typical of all the IIR satellites in that the broadcast frequency starts at a value above $+4 \cdot 10^{-12}$, decreases to a value below $+4 \cdot 10^{-12}$, and for the rest of its life is deliberately kept in the range of plus or minus $4 \cdot 10^{-12}$ of the nominal frequency. This behavior is due to the effect of TKS commands and obscures the behavior of the RAFS [3,4]. Thus, we first process the data to produce frequency curves that have the effects of the TKS removed. Figure 6 shows a curve of the RAFS frequency without TKS effects for SVN41. The slope of the RAFS frequency is the drift. Clearly, the long-term drift starts out with a large negative value and then gets smaller in magnitude, while retaining the negative sign.

We considered directly matching the drift curves, such as those shown in Figure 3 with theoretical drift curves. This seemed to be a difficult procedure, given the fluctuations we saw in the drift curves. Instead, we decided to integrate the theoretical drift curves to generate model frequency curves that we could subtract from the NGA measured frequency curves. The result was the difference in frequency between the measured data and the theoretical frequency model for an assumed drift curve. The operation of this matching procedure is described in the next section.

IV. COMPARISON OF RAFS FREQUENCY WITH THEORETICAL DRIFT MODELS

Once we had the RAFS frequency curves, the next step was to match the curves with theoretical drift models. The program **Stable** [5] allows one to match a set of frequency data with either a log fit or a diffusion fit. Equations 1 and 2 contain the formulas for the log and diffusion fits for drift and frequency, where d is the asymptotic value for the diffusion fit.

$$\text{Log drift fit} = \frac{a}{t + b} \quad (1a)$$

$$\text{Log frequency fit} = a \cdot \ln(t + b) + c \quad (1b)$$

$$\text{Diffusion drift fit} = \frac{b}{2\sqrt{t + c}} + d \quad (2a)$$

$$\text{Diffusion frequency fit} = a + b\sqrt{t + c} + d \cdot t \quad (2b)$$

Stable was used to get a best least-squares fit to the various frequency curves, which was then used to produce the difference in frequency between the fit and the RAFS data, which we called a diffusion fit residual. We found that the diffusion fit generally did better than the log fit. Even though both fits are excellent, the rest of the paper uses only diffusion fits. Figure 7 shows the results of this fitting process for SVN41 using diffusion fit. The fit reduced the frequency range of $77 \cdot 10^{-12}$ to a range of $3 \cdot 10^{-12}$. This is a very good fit, and better than we expected. Also, once we had this fit, the jumps were easily visible. As listed in Table 1, we detected 10 large jumps in the 2100 days of SVN41 initial fit data, or roughly two large steps a year. These jumps were assumed to have the standard form of a step function as described in Equation 3, where the amplitude and position were manually adjusted for a best fit.

$$y(t) = \begin{cases} 0, & t < t_0 \\ 1, & t \geq t_0 \end{cases} \quad (3)$$

We then took the frequency data, corrected for the assumed breaks, and did a second diffusion fit against the corrected data to calculate the residual, shown in Figure 8. (Note that in this step, we also corrected for steps due to velocity changes for orbit adjustments.)[6] The second diffusion fit was a much smoother curve, without the large jumps seen in Figure 7. However, we still saw little patterns after the frequency step corrections, indicating what looked to us like an overshoot process that decayed in a number of days. Thus, our model for the steps became a standard step followed by an overshoot of a number of days. As shown in Figure 9, the residual after these corrections was much smoother than the original frequency range. For some frequency steps, the magnitude of the standard step was zero and we only had an overshoot or cusp, which started with a sudden frequency step, which then decayed back to the baseline curve.

The second diffusion fit residual seemed to be composed of two sections, an initial section that decayed rapidly and then transitioned into a second section that covered the later part of the frequency curve. We attempted to model these two sections separately and found that the traditional diffusion fit was good for the second part (i.e., the frequency curve for the data after 500 days), but some satellite residuals had a

curvature. We found we could remove this curvature by assuming a diffusion fit with an asymptotic drift that was different from zero. As shown in Figure 10, the peak-to-peak frequency variation for the resulting third diffusion fit residual curve after 500 days was roughly 1 percent of the original frequency variation.

Finally, when we made these corrections, we saw that there seemed to be breaks in the curves for some satellites with what looked like drift steps in the curves. Thus, in Figure 10, which shows the residual fit for SVN41 after the above-described corrections, the residual for 500 to 800 days shows one apparent long-term drift, which is different for the period from 800 to 1800 days, which is followed by a different drift. These later breaks are smaller than the initial break, but are still significant if one estimates the drift for a number of years in advance, as is done for the IIR frequency drift corrections.

Table 1 lists the breaks that we corrected for the various SVNs using the standard step model of Equation 3. SVN41 had the greatest number of large breaks. Most of the large breaks are due to the behavior of the RAFS, but a few are related to velocity changes (“delta V”) in the satellite orbit to position the satellite in the desired location [5]. Clearly there are smaller delta V maneuvers that affected the satellite clock, but their effects were smaller than and not as obvious as those we removed in the long-term drift analysis.

The investigation into the long-term drift behavior of the RAFS has highlighted a number of items. The diffusion fit has been known as a reasonable fit for drift characteristics. However, the diffusion fit is surprisingly good for the frequency curve after 500 days or so (see Table 2). It is left for future studies to determine how much earlier one can extend the fit without suffering significant degradation. We also characterized the large frequency steps as simple step functions with an overshoot of some number of days. The initial large break in the drift characteristic has been discussed for a long time, but the smaller apparent breaks are new. The traditional drift modeling assumed that the diffusion fit would asymptotically go to zero. Our investigation indicates that zero is a good asymptotic value for most IIR clocks, but some indicate a value closer to $-1 \cdot 10^{-14}$ /day. We need more data on this issue. The diffusion fit assumes a tail of $\sqrt{1/t}$, which means that the tail decays very slowly and will be substantially away from zero for the available life of the RAFS operation with either asymptotic value. We will have to see if later data support our observation that the drift has an asymptotic value other than zero for some satellites.

V. DISCUSSION OF LARGE FREQUENCY STEPS

About half of the RAFS do not have multiple large frequency steps. Of the 13 RAFS in our study, we considered only the five SVNs 41, 45, 47, 53 and 54 as having multiple frequency steps. Of course, there are many medium-sized frequency steps that might have been considered, which would have involved most of the SVNs. However, that becomes a much more involved issue than we were prepared to handle.

We observed that the SVNs tended to have a “personality” relative to the large frequency steps. There are a variety of frequency step characteristics and the SVNs tended to have steps with selected characteristics. See the plots of the second and third diffusion fit residuals found in Table 2.

The emphasis above on the flaws of the RAFS might lead one to assume that these RAFS are performing at a poor level. On the contrary, most of the clocks have a general performance that is among the best of the constellation. Even the worst performing RAFS are near the middle of the constellation performance. Five of the 13 operational RAFS have no large frequency steps. Three others have seen only one large frequency step. That large frequency step occurred in the first year and they are currently behaving well,

with no steps. SVN44 is an interesting case, as it does not suffer from multiple large frequency steps, but it does suffer from a resonance of about a 2- to 3-day duration, which impairs its performance, so it is the worst performing operational RAFS in the GPS constellation. Of the five remaining RAFS, two have improved, so they are now in a family with high performance. Three RAFS still have substantial steps, namely 41, 47 and 53.

We noticed two general patterns of behavior that were interesting. On one hand, the RAFS seems to take some time before we see the full force of the frequency steps. For example, the first large frequency step of SVN41 in orbit was about a year after turn-on. This pattern means that the tendency of a RAFS to have large steps can go undetected in the month or two of testing at the manufacturer and before launch into orbit. A second pattern that we have seen is that the frequency and severity of frequency jumps tend to decay with time. This observation is in agreement with the general observations of the clock community. We intend to continue to monitor the constellation to see if the observed improvement lasts for the lifetime of the clocks.

VI. POSSIBLE EXPLANATIONS FOR THE OBSERVED PHENOMENA

This paper is, for the most part, an empirical observation and mathematical analysis of the RAFS behavior relative to long-term drift and large frequency jumps. This section attempts to address what we know and speculate about physical explanations for these phenomena. Camparo and associates have written papers [7,8] that identify two early intervals in the life of several types of rubidium clocks, including the RAFS. The first interval is the period where the clock reaches thermal and operational equilibrium from the turn-on transients. This interval is measured in hours and is easily observed using direct observations of the frequency output of the atomic clock. We have also observed this type of behavior, when the RAFS is under test is on the ground. However, for an on-orbit RAFS, the observations are at a distance and have to pass through the TKS system, which generates the broadcast frequency, and then have to wait until the Kalman filter has processed the monitor station observations and converged to a solution of the orbit and the clock. This convergence takes several days. Thus, by the time we get our earliest clock data, this initial interval has long since passed. Camparo and associates have identified a second interval which they call the equilibration period, which lasts for months. This period corresponds roughly to the initial sharp drop seen in our overall diffusion fits to the first break, identified above. Thus, the available theory seems to address the interval to the first break.

We identified a later interval, on the order of 2 years, where the diffusion fit matches very well. We have received some proprietary analysis from William Riley [9] describing various physical theories to describe long-term drift behaviors, including the diffusion fit. Thus, there is some theory available to address the later interval.

We have speculated as to the cause of the large frequency steps. There is anecdotal evidence that the frequency steps are related to excess rubidium, which means that some of the rubidium is in the form of liquid droplets rather than vapor. The hypothesis is that changes in the distribution of this excess rubidium can cause the observed frequency steps. What we find most interesting is that many RAFS have no steps or a few isolated steps, while other RAFS seem to repeat the same type of step peculiar to that given RAFS. We speculated that some imperfection in the physics package and an appropriate rubidium distribution leads to these steps. RAFS physics packages done “properly” don’t have these imperfections and associated steps. For the RAFS with large frequency steps, the type of imperfection and associated distribution determine the type of large step behavior. With age, the excess rubidium may be captured elsewhere or redistribute itself to produce the observed decrease in frequency steps with age. This

speculation of the cause of large RAFS frequency steps raises the thought that, if one could identify the source of the frequency steps, one might be able to remove most of the large rubidium frequency steps.

In summary, theories are available to address the items that we have observed. Hopefully, support will be available for further analysis of these issues and the manufacture of improved rubidium standards.

VII. CONCLUSIONS

The GPS IIR RAFSs provide a unique large body of data on rubidium atomic clocks. Our use of these data to examine the long-term drift and large frequency steps of the RAFS has been very successful. We were able to model the overall drift process of the RAFS with a diffusion process. This model led to an improved model, where the initial part of the curve had a decay rate faster than the diffusion model until the first break in the curve, and the later portion was an excellent fit to the diffusion model. These models were useful in determining the large frequency steps and in the detection of minor features, such as the secondary breaks and the indication that the asymptotic value of the curves was sometimes more negative than the classic assumption of a zero asymptote.

The study of the large frequency steps revealed a number of interesting results. Some frequency step components were standard step functions, but others were transients where the frequency step was followed with a decay back to the base curve. A majority of the RAFSs had zero large frequency steps or only one frequency step. Others had many steps. Each RAFS with multiple frequency steps tended to have a unique personality, leading to the speculation that the frequency step pattern was due to some imperfection in the construction of that RAFS.

The interesting results of this analysis may lead to some way of addressing these issues and provide an incentive to use this body of data to look at other RAFS phenomena.

VIII. ACKNOWLEDGMENTS

This work is based on the efforts of ITT, PerkinElmer, and Lockheed Martin, who built and now support the operation of the RAFS and the GPS IIR/IIRM satellites, and the Air Force 2SOPS (2nd Space Operations Squadron), which operates the GPS system. We also gratefully acknowledge the organizations that monitor the GPS operation and have provided us with records of clock performance, such as NGA and NRL. We also thank John Vaccaro, William Riley, and Jim Camparo who provided us with details of the RAFS operation and theory.

IX. REFERENCES

- [1] H. Rawicz, M. Epstein, and J. Rajan, 1993, "*The Time Keeping System for GPS Block IIR*," in Proceedings of the 24th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting, 1-3 December 1992, McLean, Virginia, USA (NASA Conference Publication 3218), pp. 5-16.
- [2] NGA Internet URL <http://earth-info.nga.mil/GandG/sathtml/>

- [3] M. Epstein and T. Dass, 2002, “*Management of Phase and Frequency for GPS IIR Satellites*,” in Proceedings of the 33rd Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 27-29 November 2001, Long Beach, California, USA (U.S. Naval Observatory, Washington, D.C.), pp. 481-492.
- [4] NAVSTAR Analysis Updates and Life Test Reports, Naval Research Laboratory, Washington, D.C., USA.
- [5] W. Riley, “*Stable*” Frequency Stability Analysis Program, Hamilton Technical Services, 650 Distant Island Drive, Beaufort, SC 29907, USA (<http://www.wriley.com/>).
- [6] M. Epstein, E. Stoll, and J. Fine, 2002, “*Observable Relativistic Frequency Steps Induced by GPS Orbit Changes*,” in Proceedings of the 33rd Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 27-29 November 2001, Long Beach, California, USA (U.S. Naval Observatory, Washington, D.C.), pp. 493-507.
- [7] J. Camparo, 2005, “*Does the Light Shift Drive Frequency Aging in the Rubidium Atomic Clock?*” **IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control**, UFFC-52, 1075-1078.
- [8] J. C. Camparo, C. M. Klimcak, and S. J. Herbulock, 2005, “*Frequency Equilibration in Vapor-Cell Atomic Clock*,” **IEEE Transactions on Instrumentation and Measurement**, IM-54, 1873-1880.
- [9] William J. Riley, private communication on drift mechanisms, 20 June 2006.




Figure 1. Frequency, frequency range, and age of SVN. Note that each horizontal red line shows the frequency range for a given RAFS. To the right of the line are the SVN number and the age of SVN clock. Thus, the bottom line is for SVN 41 with an age of 5.6 years.




Figure 2. Frequency range and age for each SVN.




Figure 3. Frequency drift curves for all SVNs (100-day averaging).

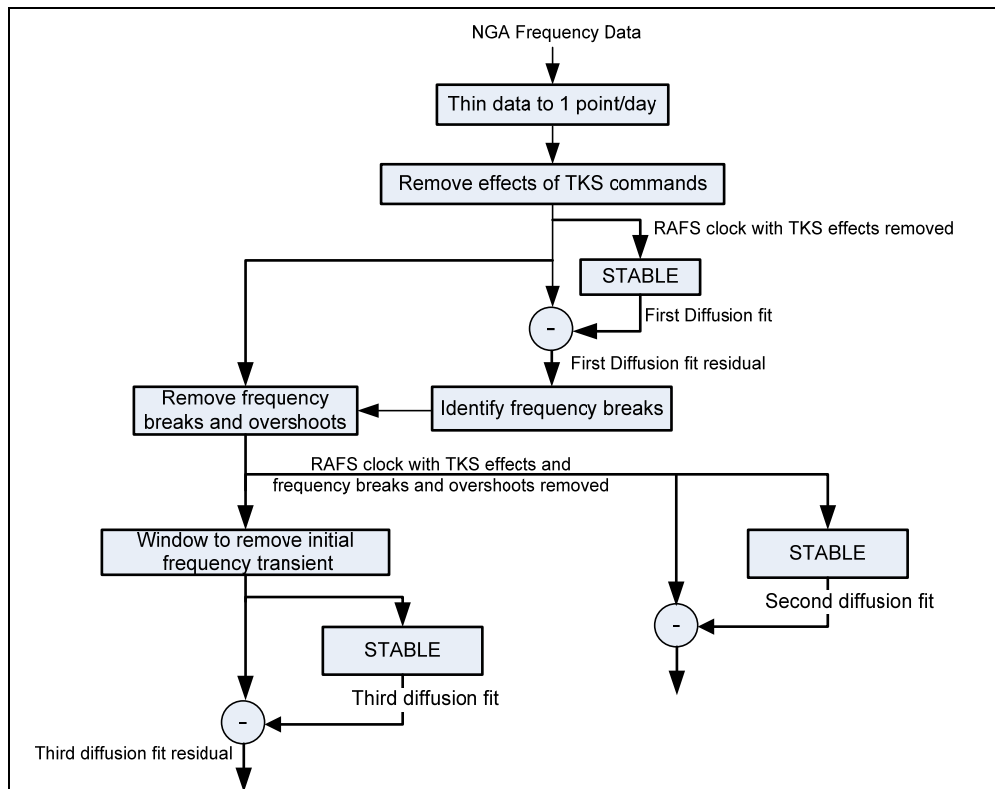


Figure 4. Processing of the NGA frequency data to reveal long-term frequency behavior.




Figure 5. Broadcast clock for SVN41 with TKS and RAFS effects present.




Figure 6. RAFS clock for SVN41 with TKS effects removed.




Figure 7. First diffusion fit residual for SVN41.




Figure 8. Second diffusion fit residual for SVN41 after deletion of large breaks.




Figure 9. Second diffusion fit residual for SVN41 after deletion of large breaks & overshoots.










Figure 10. Third diffusion fit residual for SVN41 after windowing to only include data past 500 days.






Table 1. List of frequency breaks removed from data.





SVN	Date	SV Day	Size (e-13)	Reason
41	10/28/2001	345	7.2	
	2/3/2002	445	1.7	
	3/31/2002	515	1.0	
	10/15/2002	694	7.5	
	3/5/2003	841	2.4	
	9/17/2003	1035	7.4	
	5/1/2004	1261	2.1	
	11/7/2004	1453	8.7	
	2/18/2005	1555	3.4	
	8/25/2005	1747	2.6	
44	3/16/2001	218	4.2	
46	5/13/2000	144	2.7	
47	11/13/2005	690	11.6	
53	10/16/2005	17	4.7	
	10/18/2005	19	4.0	delta V
56	11/4/2003	271	0.9	
60	7/20/2004	22	6.2	delta V
61	6/16/2005	218	1.6	delta V






Table 2. Diffusion fit residuals for all SVNs, categorized by number of frequency breaks.

SVN	Second Diffusion Fit Residuals	Third Diffusion Fit Residuals
1. SVNs with No Frequency Breaks		
43	 <p>SVN 43 Diffusion fit residuals (Stable)</p>	 <p>SVN 43 Diffusion fit residuals (Stable) (1e-14/day added to freq)</p>
SVN43 is the oldest RAFS in the constellation with high performance and no large frequency steps, so far.		


51	<p>SVN 51 Diffusion fit residuals (Stable)</p> 	<p>SVN 51 Diffusion fit residuals (Stable) 0/day removed from freq</p> 
Secondary breaks near 700 and 1650 days		
59	<p>SVN 59 Diffusion fit residuals (Stable)</p> 	<p>(not enough data beyond 500 days to make third diffusion fit residual)</p>
Excellent clock still in early life		
60	<p>SVN 60 Diffusion fit residuals (Stable)</p> 	<p>(not enough data beyond 500 days to make third diffusion fit residual)</p>
Excellent clock still in early life		

61	<p>SVN 61 Diffusion fit resid (Stable)</p> 	<p>(not enough data beyond 500 days to make third diffusion fit residual)</p>
Excellent clock still in early life		
2. SVNs with One Frequency Break		
44	<p>SVN 44 Diffusion fit resid (Stable)</p> 	<p>SVN 44 Diffusion fit resid (Stable) (1e-14/day removed from freq)</p> 
One large frequency step $4 \cdot 10^{-13}$ at 7 months. Performance suffers from 2- to 3-day resonance.		
46	<p>SVN 46 Diffusion fit resid (Stable)</p> 	<p>SVN 46 Diffusion fit resid (Stable) 0/day removed from freq</p> 
One large frequency step in its life. This step is $2.7 \cdot 10^{-13}$ in size and occurred about 5 months into the life of this 6.6-year RAFS.		

56	<p>SVN 56 Diffusion fit resid (Stable)</p> 	<p>SVN 56 Diffusion fit resid (Stable) 0/day removed from freq</p> 
<p>SVN56 has only one large frequency step in its life. This step is only $0.9 \cdot 10^{-13}$ in size and occurred about 9 months into the life of this 3.5-year RAFS. Best diffusion fit after 500 days.</p>		
<p>3. SVNs with Multiple Frequency Breaks</p>		
41	<p>SVN 41 Diffusion fit resid, cusps removed (Stable)</p> 	<p>SVN 41 Diffusion fit resid, cusps removed (Stable) (0/day added to freq)</p> 
<p>This SVN is the poster satellite for frequency steps. For a number of years, it generally had one large persistent frequency step over $7 \cdot 10^{-13}$ in size and a smaller persistent step under $3.5 \cdot 10^{-13}$ in size in each year. Also, the frequency steps generally had overshoots which quickly decayed over a number of days. This satellite didn't seem to have frequency steps in the testing at the manufacturer and the first frequency step appeared about a year after the RAFS had been turned on, in GPS orbit. SVN41 has not had a large frequency step over $4 \cdot 10^{-13}$ in almost 2 years, which has led to the speculation that SVN41 will not have any more frequency step over $4 \cdot 10^{-13}$ in size. More data are needed to see what will happen. Note secondary breaks near 750 and 1900 days.</p>		

45		
	<p>Started out with large frequency steps, which decayed in a few days. Steps spaced a week or two apart. After about 250 days frequency steps became smaller. After about 2 years, frequency steps stopped. Currently excellent clock.</p>	
47		
	<p>SVN47 has many frequency jumps, on the order of $5 \cdot 10^{-13}$, which tend to occur every week or two. Like SVN45, the jumps in SVN47 decay quickly to close to where they started. However, SVN47 is now about 2.5 years old with little improvement from its peak poor performance. Time will tell if there will be a drastic improvement in performance for SVN47, as occurred with SVN45. In addition to the many decaying frequency steps, SVN47, on 13 November 2005, had a persistent frequency step of $11 \cdot 10^{-13}$.</p>	
53		<p>(less than 500 days of data available)</p>
	<p>SVN53 is the youngest in-orbit RAFS, with about 0.8 years of life. It has a unique character compared to the other in-orbit RAFS, with a number of frequency jumps and a number of other frequency changes that seem to be unrelated to its frequency jumps. Time will tell if it will stabilize to a pattern similar to some other RAFS and if its performance will improve.</p>	

54



SVN54 has a different pattern of large frequency breaks. Frequency steps are not easily observed in the initial fast decay of the frequency. However, about 8 months into the life of SVN54, a frequency step of about $3 \cdot 10^{-13}$ was observed that quickly decayed. A number of similar frequency breaks, with each one smaller in size, were observed in the following months, until they became lost in the noise. SVN54 has another interesting anomaly about 1400 days into its life, where the frequency went down over a number of days to about $2 \cdot 10^{-13}$ below its normal value and then went back to the normal value over a number of days. We don't want to call this anomaly a frequency step, as we reserve the term frequency step to changes in frequency that are very quick, and take no more than 2 days to change. However, this anomaly, which takes many days to change, is clearly visible in the diffusion fit residuals.

